XBF: Scaling up Bloom-filter-based Source Routing

Markku Antikainen*, Liang Wang†, Dirk Trossen‡, Arjuna Sathiaseelan†
*Aalto University, †University of Cambridge, ‡InterDigital Europe, Ltd
first.last@{aalto.fi, cl.cam.ac.uk, interdigital.com}

arXiv:1602.05853v1 [cs.NI] 18 Feb 2016

ABSTRACT

A well known drawback of IP-multicast is that it requires per-group state to be stored in the routers. Bloom-filter based source-routed multicast remedies this problem by moving the state from the routers to the packets. However, a fixed sized Bloom-filter can only store a limited number of items before the false positive ratio grows too high implying scalability issues. Several proposals have tried to address these scalability issues in Bloom-filter forwarding. These proposals, however, unnecessarily increase the forwarding complexity.

In this paper, we present Extensible-Bloom-filter (XBF), a new framing and forwarding solution which effectively circumvents the aforementioned drawbacks. XBF partitions a network into sub-networks that reflect the network topology and traffic patterns, and uses a separate fixed-length Bloom-filter in each of these. We formulate this partition assignment problem into a balanced edge partitioning problem, and evaluate it with simulations on realistic topologies. Our results show that XBF scales to very large networks with minimal overhead and completely eliminate the false-positives that have plagued the traditional Bloom-filter-based forwarding protocols. It furthermore integrates with SDN environments, making it highly suitable for deployments in off-the-shelf SDN-based networks.

1. INTRODUCTION

Conventional IP routing has been suffering from state explosion at the network core. It is well known that the backbone routers have to maintain routing tables with tens of thousands of entries. This not only introduces a significant amount of management overhead, but also hinders the scalability of IP-routing and IP-multicast in particular. The issue is exacerbated by the emergence of Internet-of-Things (IoT) paradigm, where an even larger number of devices are expected to be connected over IP-enabled networks, creating even more state in the network. Furthermore, the routing operations in IP networks require advanced memory and packet switching technology in order to implement the longest-prefix matching necessary to forward packets from the sender to the receiver.

One approach to address these pinching points of IP routing is that of Bloom-filter-based source-routing[1]. In these schemes, the path from sender to receiver is computed by a dedicated path computation element, while each link along the path is identified with a domain-unique link identifier [1, 2]. Bloom filters are used to encode all link identifiers along the path from the sender to the receiver(s) within a single bitfield [1]. With that, multicast is naturally supported while the forwarding operation becomes that of a membership test in a bitfield (the Bloom filter). In other words, the forwarding decision at each switch can be realized as a simple bitwise AND/CMP operation, compared to longest prefix matching in IP routers. Applications for such simple forwarding solution can be found in multi-site VPN solutions [3, 4], in datacenters [5, 6], and in information-centric networking (ICN) architectures [7]. Within these applications, the need for centralized path computation is well aligned with the overall network architecture, while a minimal forwarding operation is specifically aimed at for performance reasons. Furthermore, work in [8] has shown that BF-based forwarding can be easily implemented in off-the-shelf SDN equipment at the cost of constant-sized flow matching tables, making it highly suitable for deployments in such emerging network environments.

However, BF forwarding comes with a significant scalability drawback. The size of the bitfield that is used for forwarding (and which holds the path information) determines the maximum number of links that can be stored on it. Thus, the size of the bitfield defines the maximum size of the network topology and the maximum size of the multicast group. Approaches presented in [1] and applications such as those presented in [3, 5, 7] usually rely on bitfields of 256 bit length, which limits their deployment capability. Extending the bitfield size is surely possible but only within the limit of defining a larger bitfield, requiring the change of all protocol headers, which is hardly a suitable approach for flexible

[1]While work on BF forwarding does not prevent its application for inter-domain routing, most solutions have been applied to intra-domain routing only. We shall therefore limit our presentation to the intra-domain problem, too.

deployment across a range of topologies.

In order to better frame the particular design choices we make for our solution to this scalability problem, we formulate the following main goals for a suitable BF forwarding solution:

1. Any solution must support any topology and group size: in contrast to existing solutions [1, 2], we are seeking a solution that does not limit the topology sizes supported.[2]
2. Any solution should allow for instantaneous multicast group formation: one aspect of BF forwarding is that multicast trees can easily be formed by unicast path information by simply bitwise ORing the unicast BF information into a single (now multicast) BF identifier. Any solution should preserve this capability in order to allow for applications such as [9].
3. The basic forwarding operation should be kept as simple as possible to the original AND/CMP based forwarding, i.e., an extensible BF scheme should not increase individual forwarding costs.
4. Any solution should be easily realizable over an SDN-based infrastructure in order to ensure easy deployment with such emerging network environments. Solutions such as [8] have outlined how to implement basic BF forwarding over SDN equipment. Any solution should preserve such capability as much as possible.

In this paper, we propose the Extensible Bloom-filter Forwarding (XBF) solution with the attempt to address the requirements outlined above. For this, we take a partitioning-based approach in which we partition the network into several parts (i.e., sub-network) so that each part contains an equal amount of links. The forwarding inside every partition happens with a standard fixed-length Bloom-filter forwarding. When a packet traverses from one partition to another, the corresponding partition-specific Bloom-filter must be loaded into the portion of the header that is used for BF matching; this copy operation causes some processing overhead. In order to reduce this processing overhead, XBF incorporates an intelligent partition algorithm which is able to exploit both network topological properties and traffic patterns, so that inter-partition traffic volume is minimised. Moreover, the resulting traffic-aware partitioning can also improve header overhead by effectively reducing the number of Bloom-filters carried in a packet. For this, we formulate the XBF design problem into a balanced edge partition problem, and present an intelligent partition algorithm that adopts the traffic-aware weighting scheme in order to minimise both header and processing overhead.

[2]In environments where we tolerate false positives in BF-based forwarding, this requirement addresses the need for supporting arbitrary topology sizes with defined false positive tolerance boundaries.

The rest of the paper is organised as follows. Section 2 presents a more detailed background on BF-based forwarding. We then describe the aforementioned zoning approach to BF forwarding as a system model in Section 3, followed by our traffic-aware partitioning approach in Section 4. In order to verify the compliance with our requirements 3 and 4, we present the experimental setup for our evaluation in Section 5. Our results regarding header and processing overhead are presented in Section 6, together with results on impact of different traffic patterns. As a proof-of-concept, specifically addressing requirements 3 and 4 again, we also present an implementation of our approach in Section 7, using a P4 software switch [10]. The paper then concludes with a discussion on possible future work.

2. BACKGROUND

2.1 Bloom-filter forwarding

Bloom-filter forwarding is an attractive scheme especially for content distribution applications due to two reasons: (1) unicast and multicast both use the same forwarding operation; (2) multicast does not require per-flow or per-group state to be stored in the network. These advantages are in stark contrast to the conventional IP-based multicast solutions where the routing table sizes grow as a function of both network size and the number of multicast groups in the network.

The basic idea behind Bloom-filter forwarding is simple. Every unidirectional link in the network is given a unique *link identifier* l, which is a m-bit long string with k independently selected random bits set to one (i.e., overlapping is possible). By so doing, any delivery path L can be described as a set of links that a packet will traverse in the network.

In general, Bloom filter routing can be split into two parts: *path construction* and *packet forwarding*. In the path construction phase, the sender or a logically centralized topology manager creates an in-packet Bloom-filter by combining the link identifiers towards every recipient together. In Figure 1, the multicast tree to v_6 and v_7 contains four links (marked as green). The in-packet Bloom-filter F is created by simply bitwise-ORing the link identifiers together, namely $F := \bigvee_{\forall l \in L} l$, where F is the in-packet Bloom-filter, L is the set of link identifiers on a path, and $\vee$ is the bitwise-OR operation Then, v_1 embeds the resulting filter `00111101` into the packet header and sends it out to the network.

The packet forwarding step becomes rather straightforward with a packet header containing all the links which need to be traversed by. Whenever a packet arrives at a switch, the switch extracts the Bloom-filter from the packet header, and for every outgoing link, it checks whether the particular link identifier has been embedded in the Bloom-filter. This operation, referred

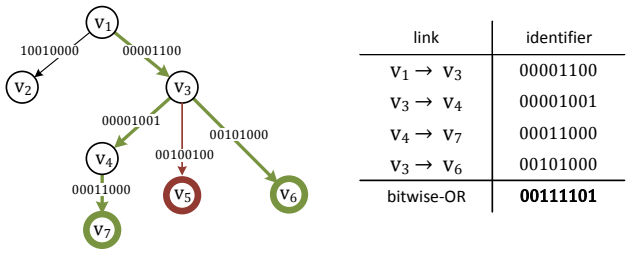

link	identifier
$v_1 \rightarrow v_3$	00001100
$v_3 \rightarrow v_4$	00001001
$v_4 \rightarrow v_7$	00011000
$v_3 \rightarrow v_6$	00101000
bitwise-OR	**00111101**

Figure 1: Bloom-filter for sending traffic from v_1 to receivers v_6 and v_7 causes a false positive and makes the packet to traverse also the link $v_3 \rightarrow v_5$.

to as *membership test*, is done by checking the equality $F \wedge l \stackrel{?}{=} l$, where l is the link identifier, and $\wedge$ denotes the bitwise-AND operation.

While Bloom-filter membership test makes the forwarding very simple, it also makes forwarding operation susceptible to false-positives. Given a fixed-length Bloom-filter, as more and more links are embedded in, the membership test may mistakenly report the presence of a link whose identifier actually has never been embedded before. Such error is specifically referred to as *false positive*. False positive indicates that a switch may forward a packet to an interface where the packet is not intended to. For example in Figure 1, a packet with Bloom-filter `00111101` will be forwarded also to v_5 even though it is only intended to v_6 and v_7. The reason is because the membership test for link $v_3 \rightarrow v_5$ produces false positive, i.e., `00111101` $\wedge$ `01001000` $=$ `01001000`.

False positive rate (FPR) is a widely used metric to evaluate a Bloom-filter-based routing scheme, which effectively tells the probability that a switch forwards packet to an unintended interface. FPR is a function of the Bloom-filter parameters (i.e., values of m and k) and the number of link identifiers stored in the filter. It is obvious that FPR will be high for large multicast trees, since more links are involved in delivery and their identifiers have to be embedded. High FPR values can cause a significant amount of redundant traffic increasing the probability of congestion in the network. Another problem with false positives is that they may cause a packet to be forwarded back to previously visited nodes in the multicast tree and thus render (possibly infinite) forwarding loops.

We now continue by describing how earlier research has approached these challenges.

2.2 Challenges of BF-forwarding

The scalability-issues caused by the fixed Bloom-filter length are well known and have been identified already in early proposals such as Free-Riding Multicast [11] and LIPSIN [1]. In general, there are three approaches that can be found from the literature regarding increasing the scalability of Bloom-filter source routing. The first, and simplest way to achieve scalability is to disallow too large multicast-trees and instead use several trees when serving content to a very large audience [12, 13]. While the solution is very simple and effective, it naturally increases the amount of traffic in the network and is suitable only in fairly small networks.

The second way to improve the scalability of Bloom-filter forwarding is to use some method to decrease the FPR. The idea is that when the FPR gets smaller, a fixed sized Bloom-filter can store larger multicast trees. One commonly used way to reduce FPR is to use several coexisting link-identifiers for every link: when the multicast trees are constructed, the link identifiers are chosen so that the number of false positives is minimized [1, 14]. The drawback of this approach is it increases the amount of state stored in the switches. Also, these approaches only slightly decrease the FPR and do not provide scalability for arbitrary large networks. Furthermore, because the FPR still remains positive, these forwarding schemes must employ extra countermeasures against forwarding loops caused by false-positives thus further complicating the router design [15, 16].

The third way to solve the scalability problem is to encode the multicast-trees into variable sized Bloom-filters [15, 17]. This basically means, that when a multicast tree grows too large (and thus causes too may false positives) the in-packet Bloom-filter is dynamically extended to keep the FPR acceptable. This requires that there is an additional length-field in the packet header which tells the length of the packet's Bloom-filter. These variable-length Bloom-filter solutions would appear palatable due to the fact that it scales to indefinitely large networks and multicast groups.

Making the Bloom-filters variable-length has, however, two major drawbacks. First, it makes the multicast-group management more difficult. This is because variable-length Bloom-filters cannot be combined together simply by bitwise ORing them together (they are of different lengths, after all). Because of this, the topology manager becomes a bottleneck as it must be consulted whenever the multicast tree changes (i.e., a node joins or leaves the multicast group). The second problem is that variable-length Bloom-filters vastly increase the complexity of the forwarding. When the Bloom-filter lengths vary from packet to packet, the link identifiers have to be calculated based on the length of the Bloom filter. Thus, the switches have to calculate link-identifiers at line-speed. Alternatively, the switches have to store pre-calculated link-identifiers for every possible Bloom-filter length. This, on the other hand, would significantly increase the amount of state stored in the network.

There are also proposals that try to combine the aforementioned approaches. For example, Multi-stage Bloom-

filter (MSBF) tries to reduce the false positives by splitting the multicast-tree into several variable-length stages which use different Bloom-filters [2, 18, 17]. The problem with this kind of solutions is, again, that they tend to make group management very difficult and increase the complexity of the forwarding nodes.

3. XBF FORWARDING SCHEME

In this section, we introduce the extensible Bloom-filter forwarding (XBF) and the system model under which it operates. XBF addresses the fundamental scalability issue of conventional Bloom-filter forwarding by partitioning a network into equisized parts inside which a separate Bloom filter is constructed.

3.1 System Model

XBF forwarding scheme is designed for wide scale intra-domain multicast. The scheme works in networks where roughly half of the switches can perform packet matching on arbitrary fields (e.g. software switches or P4-switches) while the rest of the switches can be regular OpenFlow-switches.

We assume a network that can be represented as a directed graph, $G_n = (V_n, E_n)$, where V_n is the set of switches and E_n is the set of links[3]. The basic idea in XBF is to group the network edges into partitions so that every partition contains at most 256 edges. Then, every partition is allocated with a separate 256-bit partition Bloom-filter. Because the partition Bloom-filters are 256-bit long and every partition contains at most 256 edges, we can assign a *unique one-bit link identifier* for every link instead of creating the link identifiers randomly. Therefore we can completely get rid of false-positives.[4] The maximum partition size is chosen as 256 because it fits into the source and destination address fields of a IPv6 header.

Figure 2 presents the XBF packet format. The packet header has two parts: *iBF* contains the Bloom-filter used for intra-partition forwarding, and *zBF* contains the Bloom-filters of all of the partitions which the packet needs to traverse by. The content in zBF remains the same throughout a delivery whereas the content in iBF changes dynamically depending on which partition the packet is currently in. Together the iBF and zBF constitute the header overhead.

There are three main entities in a XBF-enabled network: *topology manager*, *popper switches*, and *forwarder switches*. The topology manager is a logically centralised entity similar to a SDN controller, with three main functions: (1) it partitions the network so that each partition contains at most 256 directed links; (2) it assigns the one-bit link identifiers for every link in the network; (3) it creates Bloom-filters that can be used to send packets from a source to a destination. Note that (1) and (2) are only performed once for the whole system while (3) is performed once for a new multicast tree.

[3]In this paper, we use the terms *node* and *switch* interchangeably as well as the terms *edge* and *link*.

[4]Technically, the described data-structure is no longer a Bloom-filter. We nevertheless abuse this term as the operations done with these one-bit identifiers match the Bloom-filter operations.

F	in-packet Bloom-filter, $F \in \{0,1\}^m$
m	length of a Bloom filter
$\wedge$, $\vee$	bitwise-AND and bitwise-OR respectively
$G = (V, E)$	directed graph with vertices V and edges E
$G_n = (V_n, E_n)$	directed network graph
$G_c = (V_c, E_c)$	directed connectivity graph, $V_c = E_n$
$\tau_{a,b}$	traffic volume on network link $v_a \rightarrow v_b$
ν	imbalance coefficient in vertex partitioning
P	set of partitions
p_e	partition to which a link (or connectivity graph vertex) e belongs to, $p_e \in P$
$N(v)$	partitions that are adjacent to node v, $N(v) \subseteq P$
ϕ_e	number of popping operations done on each packet traversing link e

Table 1: Notation used in this paper

A popper switch (or popper) refers to a switch sitting at the boundary of partitions (e.g., v_1, v_2, and v_3 in Figure 2). More specifically, a switch v is a popper iff the edges incident to v belong to at least two partitions. Meanwhile, a forwarder switch (or forwarder) refers to a regular switch with all its links allocated to the same partition. Forwarder switch is a simple device that only performs basic Bloom-filter forwarding operations – it checks which of its next-hop links are included in the iBF then forwards the packets to the corresponding interface. This simple operation, however, only applies to the forwarding occurs inside one partition. If a packet is intended to another partition, it must go through a popper which lies at the border of multiple partitions. Unlike the basic forwarding, a popper copies a partition Bloom-filter from the zBF to the iBF every time it performs the membership tests for the links of the partition. We refer this operation as *popping* in the paper.

Figure 2 illustrates popping operations. The multicast tree (marked as orange arrows) crosses partition borders at two different popping nodes: v_1 and v_3. When the packet arrives at v_1, the popper has to pop (i.e., copy) the blue Bloom-filter from zBF to iBF. After popping, v_1 can check to which outgoing blue interfaces the packet should be forwarded as a regular forwarding element. However, the operation is slightly different at popper v_3 who has to pop the packet twice. The reason is that v_3 sits at the border of three partitions hence it needs to perform the popping for both green Bloom-filter (partition 3) and blue Bloom-filter (partition 2). As we can see, the number of popping operations a popper needs to perform for a packet depends on the number of incident partitions.

3.2 Optimising Popping Via Edge Partition

From both network operator's and end user's perspec-

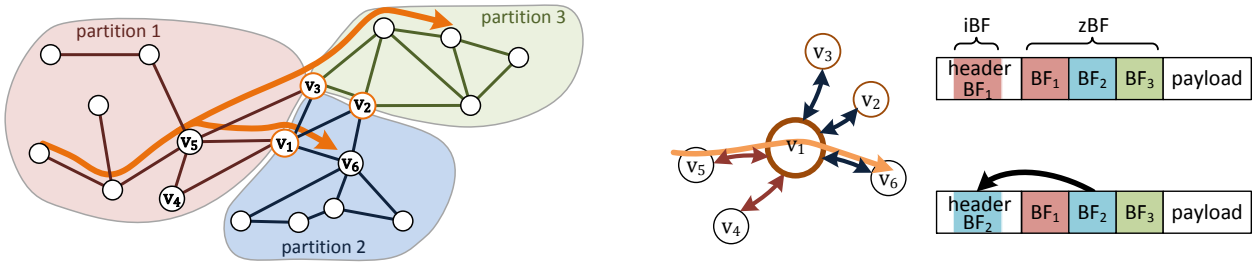

(a) A sample topology with three partitions, popping switches are v_1, v_2, and v_3 at the boundary of partitions.

(b) The packet header format and the popping operation at v_1 when the packet traverses the path $v_5 \rightarrow v_1 \rightarrow v_6$.

Figure 2: XBF system model, a naive multicast tree is used to illustrate header format and popping operation.

tive, the total number of poppings in the network needs to be minimised, which requires the topology manager to strategically select poppers to connect different partitions. To help us in formulating the optimisation problem in Section. 4, we first define the required number of popping operations as follows. Let p_e denote the partition which a directed link e belongs to. Let H_e be the set of possible next-hop links that the packets can traverse after link e. The following equation gives the number of poppings required for each packet traversing via e (note that the set cardinality tells the number of *distinct* elements in it).

$$\phi_e = \left| \left(\bigcup_{e' \in H_e} \{p_{e'}\} \right) \setminus \{p_e\} \right| \quad (1)$$

There are two technical details worth mentioning: First, the set of possible next-hop links ϕ_e excludes the interface where a packet originally came from in order to avoid loops. Second, the popping operation is not necessarily needed for all passing-by packets. A popper only performs popping for the packets entering into different partitions, i.e., cross-partition traffic. For the packets remaining in the same partition after forwarding, popping operation can be skipped. We will present the actual implementation details in Section 7.

4. TRAFFIC-AWARE PARTITIONING

We have designed Jigsaw, an algorithm to calculate the optimal partitioning of a given network to minimise poppings. Jigsaw is traffic-aware and is also an important component in topology manager, playing a significant role in XBF performance.

In brief, Jigsaw aims to minimise the traffic volume across partition boundaries (i.e., inter-partition traffic) in order to minimise the total number of poppings. More specifically, for a given traffic pattern and a network topology, Jigsaw optimises the partitions boundaries so that majority of the traffic occurs within a single partition, while simultaneously keeping the number of links in each partition constant (or near-constant).

There are three steps in the algorithm:

1. **Weighting**: the algorithm adds weight to every link in the network based on the actual traffic volumes.
2. **Translation**: the algorithm transforms the original network graph into a *connectivity graph*.
3. **Partitioning**: the algorithm performs balanced vertex partitioning on the connectivity graph to obtain an optimal partitioning.

We now separately describe the details of these steps.

4.1 Embedding weights in original network

The very first step in Jigsaw is assigning weight on every directed link according to the traffic volume passing by. Since forwarding and popping operations are performed on per-packet basis, we set the weight of a link equal to the number of packets passing through in a well-defined time-unit (e.g., an hour).

Although weighting indicates that a topology manager needs to possess the global knowledge of traffic volume in a network, if such information is not available, the traffic volumes can be approximated with switches' betweenness centrality values. Many prior work have shown that there is a strong correlation between traffic patterns and networks' topological properties [19, 20]

4.2 Constructing a connectivity graph

The second step is to construct a connectivity graph [21] from the original network. The algorithm for creating such a connectivity graph is simple. Firstly, for each directed *link* (v_i, v_j) in the original graph, a matching *vertex* (v_i, v_j) is added to the connectivity graph. Accordingly, the original link weight becomes the vertex weight in the connectivity graph. Then edges are added to the connectivity graph with a simple rule: a directed edge is added between two connectivity graph vertices iff a packet can traverse from the corresponding network link to another in the original graph. Figure 3 illustrates a connectivity graph created from a simple network graph.

There are several motivations of constructing a con-

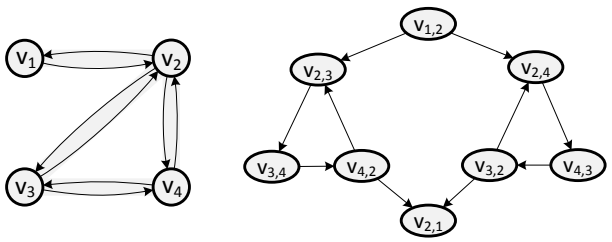

Figure 3: A network graph and a connectivity graph created from it. Link weights (network graph) and vertex weights (connectivity graph) are omitted.

nectivity graph. First, we can translate a balanced edge partition problem to an equivalent balanced vertex partition problem. While the two problems have been shown to be equivalent [22], balanced vertex partitioning is already well studied hence there are many mature algorithms[5].

Second, a connectivity graph is able to reflect actual paths more accurately and further remove many potential loops. For example in Figure 3 where a packet traverses from v_3 to v_2, it is obvious that the next-hop link cannot be $v_2 \rightarrow v_3$ since that is where the packet comes from. The directed network graph does not encompass this information about physical interfaces, whereas connectivity graph does not have this limitation – an edge between two connectivity graph vertices means by definition that a packet can follow that path.

Third, a connectivity graph is able to take traffic direction into account, which can lead to better partitioning especially when traffic is asymmetric in different directions. This can be seen by examining the connectivity graph in Figure 3: the connectivity graph seems to contain two parts that are very weakly connected to each other (i.e. $v_{2,3} \rightarrow v_{3,4} \rightarrow v_{4,2}$ and $v_{2,4} \rightarrow v_{4,3} \rightarrow v_{3,2}$). This hints that an optimal partitioning may, rather counter-intuitively, may place reverse links to different partitions. This is also what we observed when partitioning realistic network graphs (see Section 6). An edge partitoining done directly on the network graph could not

4.3 Applying vertex partitioning

By definition, the balanced edge partitioning problem in the original graph can be naturally solved by applying a balanced *vertex* partitioning on the corresponding connectivity graph. The goal is to group the vertices of a given connectivity graph $G_c = (V_c, E_c)$ into n components of at most 256 vertices, while trying to minimise a well-defined *total cost* function.

We use τ_v to denote the traffic volume on a connectivity graph vertex $v \in V_c$ (remember that every connectivity graph vertex uniquely describes a link in the physical network). Also, we use $p_v \in P$ to denote the partitions of v, and $N(v) \subseteq P$ to denote the set of partitions to which v's neighbours belong to. The total cost of a partitioning is now defined as follows:

$$totalv = \sum_{v \in V} \tau_v \cdot |N(v) \setminus \{p_v\}| \quad (2)$$

Comparing to Eq. 1, besides the fact that traffic volume has been taken into account, the only difference in Eq. 2 is that the calculations are performed over all vertices in the connectivity graph (corresponding to the links in the original graph). Therefore, it is not difficult to realise that *totalv* equals to the total number of popping operations in the network. Consequently, the optimal partition can be obtained by finding a solution that minimizes the *totalv*:

$$\arg\min_{\forall v \in V, p_v} totalv \quad (3)$$

This optimization problem is well known and is called as graph partitioning while minimizing the total communication volume [23]. The problem has been shown as NP-hard. However, there are efficient approximation algorithms that can solve the problem in polynomial time given the condition for equisized partitions is slightly relaxed [24]. In our case, since we only require that the partition sizes *do not exceed* 256, we can simply apply, for example, a multilevel k-way partitioning algorithm to find a nearly-optimal solution for the problem [25]. This algorithm is part of widely used vertex partitioning library called METIS [26].

Finally, because of the one-to-one mapping between a connectivity graph and a network graph, the vertex

Data: Network graph $G_n = (V_n, E_n)$. Every $e \in E_n$ is a 3-tuple $(v_a, v_b, \tau_{a,b})$, where v_a and v_b denote the link start and endpoints respectively, and $\tau_{a,b}$ denotes the traffic volume on that link.
Result: Partitioned graph

```
// Create connectivity graph and annotate it with
   link weights
foreach (v_a, v_b, τ_{a,b}) ∈ E_n do
    V_c := V_c ∪ {(v_a, v_b, τ_{a,b})};
    foreach v_c ∈ getNeightbours(G_n, v_b) do
        if v_a ≠ v_c then
            E_c := E_c ∪ {((v_a, v_b, τ_{a,b}), (v_b, v_c, τ_{b,c}))};
        end
    end
end
// Apply the vertex-partitioning algorithm
G_c := (V_c, E_c);
no_partitions = (|E|/256) × ν;
return balancedVertexPartitioning(G_c, no_partitions)
```

Algorithm 1: Jigsaw, the traffic-aware network partitioning algorithm

[5]To be more specific, the edge-partition problem with *no aggregation* is equivalent to a balanced vertex-partition problem. This corresponds to the problem we are solving.

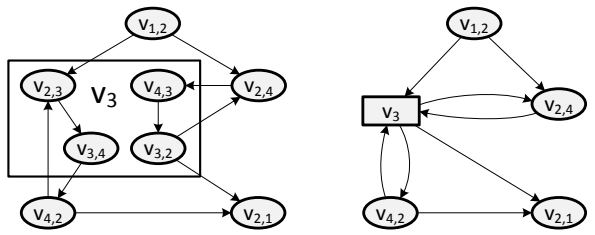

Figure 4: Collapsing the connectivity-graph vertices adjacent to v_3 switch if it cannot act as a popper switch.

partitioning obtained on the connectivity graph can be easily translated back to the edge partitioning of the original graph. Algorithm 1 presents the key operations in Jigsaw algorithm.

4.4 Partitioning heterogeneous networks

So far, we have assumed that any switch can be selected as a popper. In some networks, however, only a subset of switches can perform the popping function while the rest, referred to as *legacy switches*, can only act as forwarders. For example, regular OpenFlow switches cannot perform popping but can forward packets based on the iBF found in the IPv6-address fields [27]. We now briefly explain how the Jigsaw algorithm can be easily extended to handle these situations.

Recall that a switch is a popper iff it has incident links belonging to two or more different partitions. Thus, a legacy switch (i.e. forwarder) must have all of its links belonging to a single partition. This constraint can be fulfilled by collapsing all connectivity graph vertices together which represent links incident to a legacy switch. Figure 4 illustrates this in a scenario where v_3 is a legacy switch. All links incident to switch v_3 are aggregated to a single connectivity graph vertex. In this way, we do not have to change the partitioning algorithm at all, but simply count vertex v_3 as four vertices during the partitioning phase of Jigsaw.

Due to space constraints, the collapsing operation has been omitted in Algorithm 1. Moreover, for the sake of clarity, in the rest of this paper we consider only networks where all switches can act as popper switches.

5. EVALUATION METHODOLOGY

In this section, we explain our experiment setup and the evaluation metrics.

5.1 Experiment setup

We use three different kinds of topologies in our evaluations: eight conventional ISPs, one community network, and various synthetic topologies. The ISP backbone topologies are based on the Rocketfuel r_0 datasets [29]. We dropped the smallest ISP networks having less than 256 links. For the community network, we used the Guifi.net core network in Catalonia region [30]. The synthetic topologies are generated using Barabási–Albert (BA) and Erdős–Rényi (ER) models with various parameters (shown in Table 2). For all aforementioned topologies, we only consider the largest component.

Algorithm	**Parameters**
Barabási–Albert $G(n,m)$	$n \in \{500, 1000, \ldots, 3000, 4000, \ldots 8000\}$ $m = 2$
Erdős–Rényi $G(n,p)$	$n \in \{500, 1000, \ldots, 3000, 4000, \ldots 8000\}$ $p = (1+\epsilon) \cdot n^{-1} \cdot \ln n$, $\epsilon = 0.1$

Partitioning algorithm	**Parameters**
Jigsaw	#partitions $= 1.1 \cdot (\lvert E \rvert / 256)$ Vertex-partitioning with METIS-library [26] METIS parameters: ptype: k-way, iptype: grow, ncuts: 1, objtype: vol
Powergraph [28]	target # links in a partition: 256 algorithm applied recursively for partitions with more links

Table 2: Parameters used to create synthetic topologies and in partitioning algorithms

For traffic pattern, we use both realistic (for Guifi network) and synthetic traces (for ISP and synthetic networks). The realistic trace was obtained with a 30-day monitoring of Guifi network. To generate synthetic traffic, we let every node communicate with any others with equal probability, namely we randomly select a certain number of data sinks for a given source with uniform distribution. A multicast tree is further constructed by combining the shortest paths between the data source and sinks. Both realistic and synthetic traces contain the information of traffic volume on every directional link. Figure 5 shows the traffic distribution over all the links. Both trace types exhibit similar statistical distributions which justifies our traffic generation method. We assume that the topology manager has *a priori* knowledge about the traffic distributions in the network.

The vertex partitioning (i.e. the third step in our Jigsaw-algorithm, see Section 4.3) was implemented with METIS-library [26], which we observed to produce high quality partitionings very fast (it took less than 20 seconds to partition the largest topologies with a mid-range laptop using one processor core). To evaluate the overall effectiveness of Jigsaw, we chose a popular edge partitioning heuristic called Powergraph [28] as our comparison baseline. The simulation was implemented with Python 2.7 using the Networkx, FNSS, pyMetis, and Scipy libraries. For each configuration, we repeated the experiments 1000 times to obtain statistically valid and representative results.

5.2 Metrics

In addition to the obvious metric like total number of poppers, we also calculate the following metrics to

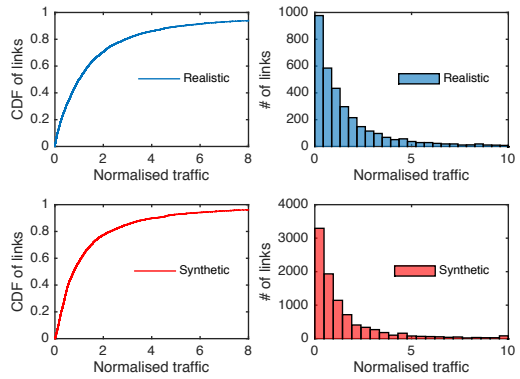

Figure 5: Comparison of synthetic and realistic traffic patterns. Realistic traffic pattern was obtained with a 30-day monitoring of Guifi-network (Catalonia-region).

evaluate the effectiveness of XBF scheme.

- Mean number of partitions through which multicast trees traverse. This can be used to further calculate the header overhead of XBF.
- Mean number of poppers through which multicast trees traverse, and mean number of popping operations for a single multicast packet. These can be used to evaluate the processing overhead in forwarding.

As mentioned, these metrics can be used to evaluate both the header overhead and the performance overhead of XBF. Moreover, to gain a comprehensive understanding of header overhead, for each topology, we calculate the minimum bit-lengths for Bloom-filters that would be required to store a multicast tree while keeping the number of false positives under a certain threshold. More specifically, we calculate values for function $L_p(s)$, which tells the minimum bit-length for a Bloom-filter that can store a multicast tree with s sinks in a given topology, while keeping p percentage of the multicast trees free from false positives. Thus, $L_{99\%}(10)$ denotes the length of a Bloom-filter that, given a topology, can store a multicast tree with 10 sinks while keeping 99% of the multicast trees free from false positives.

6. EVALUATION RESULTS

The most important evaluation results have been compiled into three groups (i.e., unicast with just one sink, multicast with ten sinks and with 20 sinks) in Table 3 along with some statistics about the topologies. Throughout this section, the whisker bars in the box-plots represent the 5th and 95th percentile values of the measurement while the boxes show the interquartile range.

6.1 Header overhead

We start by analysing the header size in XBF. Recall from Section 3 that the header overhead in XBF is the aggregated size of iBF and zBF. The size of the iBF is always 256 bits, while the length of the zBF varies depending on the number of filters in it, which again depends on the number of partitions a multicast tree traverses through. Also, since the zBF is accessed only by relatively few nodes (that is, the popper switches), it may be compressed. Here, we present the header overheads with and without such compression. The compression was done simply by compressing the entire zBF with with binary run length encoding (we use *Elias-gamma coding* for the run lengths [31]).

Table 3 shows the mean header overheads in all topologies when the multicast tree has 1 (unicast), 10, or 20 sinks. It can be seen that the header overhead is largely a function of the multicast tree size – although the topology also affects it, its effect is relatively moderate. This is further illustrated in Figure 6 which plots the header overhead as a function of different tree sizes in AS 1221 topology[6]. The horizontal dashed lines indicate the optimal bit-length for a Bloom-filters that can store a multicast tree with s sinks while keeping 99% of multicast trees false-positive free. More formally, they show the value $L_{99\%}(s)$ for $s \in \{10, 30, 60\}$.

It can be seen from the figure that when a multicast trees grows, the size of the non-compressed XBF header increases accordingly since more partitions are involved in a distribution. However, zBF can be easily compressed due to its sparsity. In other words, using a single bit to identify one link will create multicast tree representations mostly filled with zeroes. (This is not necessarily true for traditional LIPSIN-like Bloom-filters forwarding where several pseudo randomly selected bits are set to one for every link that is included in the multicast). Thus, with the simple run-length encoding done on the the zBF, it is possible to achieve up to 70% saving on header overhead.

6.2 Processing overhead in forwarding

The key goal of XBF is to allow false-positive free forwarding while simultaneously minimizing the processing overhead of the packet forwarding. While basic forwarder switches can forward packets with minimal overhead, the same is not true for popper switches. We now evaluate the *processing overhead* that is caused by the inter-partition traffic and popping operations. The key metrics used here are (1) the number of popper switches in the network, (2) the number of popper switches on a multicast tree, and (3) the number of popping operations occurring during packet forwarding.

As already discussed, the processing overhead is dependent on how the network is partitioned. Thus, the performance of Jigsaw algorithm plays a central role in the performance of XBF. We now compare Jigsaw algorithm to another relatively simple edge-partitioning

[6]There is no particular reason why we created the figures using the data from AS 1221.

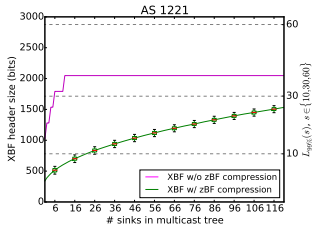

Figure 6: XBF header overhead. For the horizontal lines, see text

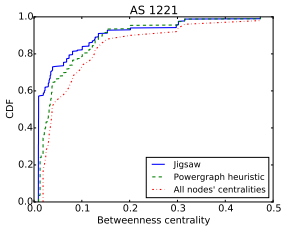

Figure 7: Betweenness centralities of the 50 most central popper nodes

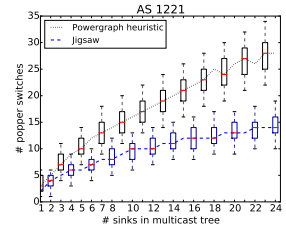

(a) Poppers on a multicast tree as function of no. sinks

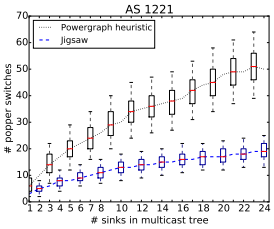

(b) No. popping operations as function of no. sinks

Figure 8: Comparison of partitioning algorithms

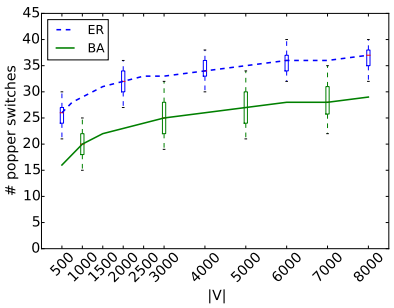

(a) No. popper switches through which a multicast tree (10 sinks) passes

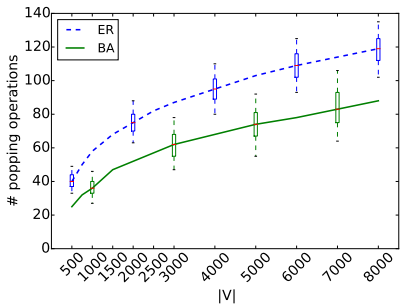

(b) No. popping operations that occur during a multicast transmission

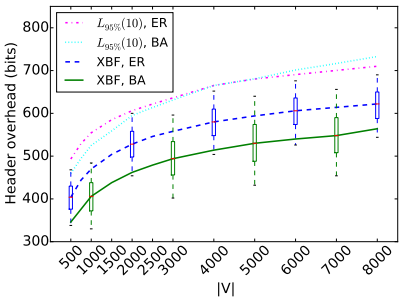

(c) XBF header overhead (with compression) and $L_{95\%}(10)$

Figure 9: Different metrics calculated from a multicast tree with 10 sinks, as a function of topology size (number of nodes). Some error bars omitted for clarity.

heuristic called Powergraph. The purpose of this comparison is to validate the performance of Jigsaw and to show the importance of the partitioning algorithm.

We observed that our Jigsaw-algorithm outperforms the Powergraph-heuristic consistently in all topologies. A notable difference between the performances of Jigsaw and Powergraph heuristic is that the partitioning given by the Powergraph heuristic requires much more popping switches. Furthermore, Powergraph-heuristic places the popper switches to more central locations (i.e., network core) compared to the Jigsaw-algorithm. This can be seen in Figure 7 which shows the betweenness centrality values of the most central popper nodes – the popper switches' centrality values are much higher when the partitioning is done with Powergraph heuristic compared to the partitioning done by Jigsaw. This basically means that the Powergraph heuristic not only produces much more popping nodes, but also places them to worse places in the network.

The placement of the poppers directly affects the number of popping operations that must be performed in packet forwarding. As a result, it is not surprising that the number of popping operations is also smaller when the network is partitioned using Jigsaw. This can be seen in figures 8(a)–8(b) which show the number of poppers on a multicast trees as well as the number of popping operations that happen during forwarding. Furthermore, the figures show that the consequences of a poor partitioning are emphasized when the multicast trees grow.

6.3 Scalability of XBF

We used the Barabási–Albert (BA) and Erdős–Rényi (ER) models to evaluate how XBF scales as the network grows. The parameters we used with these models are summarized in Table 2.

Figure 9(a) shows the number of popper switches through which a packet traverses as a function of network size. The growth in is clearly sub-linear. The sub-linearity comes largely from the fact that the network diameter grows sub-linearly – the number of partitions on a path is limited by the network diameter. This guarantees the scalability of XBF in a very natural way.

The same sub-linear shape can also be seen in the number of popping operations (Figure 9(b)) and in the header overhead (Figure 9(c)). However, the number of popping-operations (which approximates the processing overhead of the forwarding) grows slightly faster than the number of popper-switches. This is caused by the fact that as the network grows, the probability of a popper-switch having adjacent links belonging to more than two partitions grows. Nevertheless, both the growth in the popping operations and the growth of

Topology		# nodes	# directed links	Mean path len. (nodes)	# partitions	# popper switches	Mean # popping switches on path	One sink (unicast)				Ten sinks (multicast)				20 sinks (multicast)			
								Mean header overhead	Mean hdr. overhead (compressed zBF)	$L_{95\%}(1)$	$L_{99\%}(1)$	Mean header overhead	Mean hdr. overhead (compressed zBF)	$L_{95\%}(10)$	$L_{99\%}(10)$	Mean header overhead	Mean hdr. overhead (compressed zBF)	$L_{95\%}(20)$	$L_{99\%}(20)$
Rocketfuel AS	1755	172	762	5.9	4	19	1.5	758	338	83	103	1251	583	514	617	1279	725	849	1013
	3257	240	808	6.6	4	35	1.4	772	346	90	112	1244	583	496	595	1276	726	837	999
	3967	201	868	5.8	4	21	1.4	762	337	79	98	1256	592	519	622	1279	733	868	1302
	1221	318	1516	6.0	7	72	2.3	945	346	96	117	1870	602	510	765	2013	761	847	1270
	7018	631	4156	6.0	18	191	2.9	992	347	90	111	2884	706	650	975	3691	953	1128	1693
	1239	604	4536	5.2	20	310	3.1	982	336	79	96	3168	683	573	860	4133	939	1031	1546
	2914	960	5642	6.2	25	308	3.4	1078	352	96	117	3613	762	690	1035	4861	1079	1242	1864
	3356	624	10596	4.4	46	461	2.8	929	325	73	88	3865	655	535	803	5768	933	1000	1501
Barabási–Albert		500	1992	4.7	9	287	2.8	875	328	70	85	2173	599	459	689	2453	793	817	1225
		1000	3992	5.0	18	646	3.1	961	334	79	96	3107	662	519	779	3994	916	948	1422
		2000	7992	5.3	35	1432	3.6	1031	340	87	105	4020	722	592	887	5739	1035	1076	1615
Erdős–Rényi		497	3496	4.4	16	468	3.1	854	324	57	71	2875	660	491	588	3661	909	930	1395
		999	7732	4.6	34	979	3.5	931	328	62	77	3956	724	552	658	5648	1044	1051	1577
		2000	16730	4.8	72	1970	3.8	990	333	66	83	4883	783	607	910	7649	1173	1174	1760
Guifi, Catalonia		7168	14872	11.7	64	1145	5.3	1617	433	222	333	5779	1042	1190	1785	8227	1468	1946	2919

Table 3: Summary of the evaluation results on both realistic and synthetic topologies using different number of sinks.

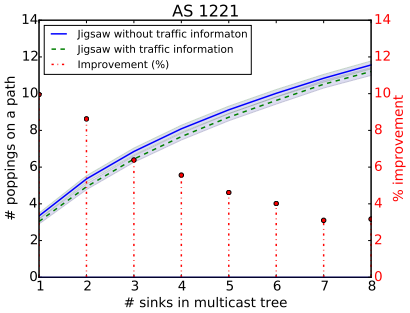

Figure 10: Partitioning quality improvement when traffic patterns are taken into account (details in text)

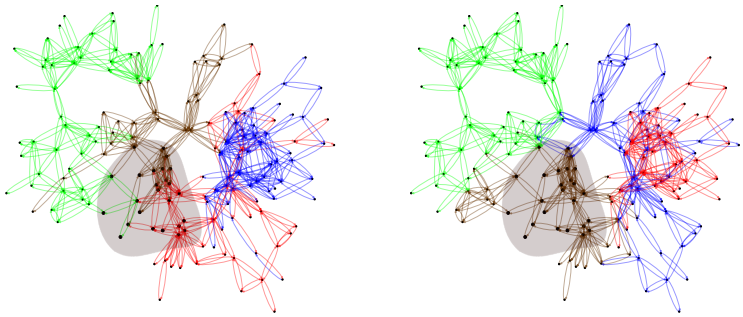

Figure 11: AS 1755 topology partitioned into four using uniform traffic distribution (left) and non-uniform traffic distribution (right). Details in text.

the header overhead are still sub-linear implying good scalability.

6.4 Effect of the traffic patterns

So far in the evaluation that every node sends equal amount of traffic to every other node. In order to gain understanding on how non-uniform traffic volumes may change the partitioning quality, we also ran the simulation with non-uniform traffic volumes. This simulation was executed as described in Section 5.1 with the difference that 10% of the nodes were assigned as high-demand nodes, which produced ten times of the traffic that a normal node does.

The effect of this non-uniform traffic in AS 1221 topology can be seen in Figure 10, which shows the mean number of popping-operations when the partitioning was done with Jigsaw-algorithm with and without taking the traffic volumes into account (the shaded areas indicate the standard deviation). The red stem-lines indicate the improvement when the traffic volumes were taken into account. As we suspected, the partitioning quality can be improved relatively much when the traffic patterns are non-uniform – a typical improvement in this experiment was roughly 9% in terms of the number of popping-operations when the traffic volumes are taken into account.

It should be noted that we did not incorporate any kind of spatial preference to our simulation (i.e. a node would be more likely to communicate with nearby nodes). To illustrate what happens when nodes communicate mostly with nearby nodes, we repeated the experiment in AS 1755 topology so that the high-demand nodes were placed near to each other. Figure 11 shows how the partitioning changes in this situation: when the traffic volumes are taken into account the partitioning algorithm places the high-demand nodes (under the shaded

area) to the same partition.

Based on these examples, we believe it is safe to say that results obtained assuming uniform traffic volumes show only the worst case performance of XBF and Jigsaw.

6.5 Summary of the results

To summarise, XBF has been shown to be able to effectively reduce both header overhead and processing overhead in our evaluation. For conventional BF-based routing, both overhead are a function of total number of links which can grow fast in large networks. Whereas for XBF, by partitioning a network, such overhead are bounded by the network diameter (especially for small multicast trees) which grows much slower due to small-world effect. The reason is because a single packet will only traverse a small amount of partitions in the network. Although as the multicast tree grows, XBF eventually includes all the partitions in zBF, it still has significantly smaller overhead comparing to conventional BF-based routing. Moreover, XBF is able to exploit traffic pattern when applying Jigsaw to further optimise its performance. Traffic awareness guarantees the volume of cross-partition traffic is minimized.

7. SDN INTEGRATION

So far, we have been focusing on describing the overall XBF architecture and understanding its system behaviours. However, for a practical deployment, it is equally important to understand the complexity and flexibility in the realistic system implementation. Thus, we implemented the XBF popper and forwarder in P4 and verified the correctness of the scheme with emulations using the *bmv2* P4 software switch [32]. We chose P4 as the implementation language since it is currently one of the most promising SDN technology allowing packet processing on arbitrary fields. Our P4 implementation also serves as a sanity check that the packet processing does not include any operations that would be difficult to implement with hardware. The P4 implementation is about 200 lines of code indicating its low complexity.

In our implementation, the XBF header is within a IPv6 header, which contains the iBF in the source and destination address fields. The XBF header itself contains two main parts: a *filter bitmap* and a *partition Bloom-filter set*. The filter bitmap has a length of $|P|$ (number of partitions in the network) and tells which partition filters are included in the partition Bloom-filter set zBF (the bitmap has a 1-bit in position n iff the multicast tree passes the partition n). The partition Bloom-filters are included in the XBF header in ascending order based on the partition identifier selected from $\{1, \ldots, |P|\}$. This makes it is possible to calculate the offset of any partition Bloom-filter using the information found from the filter bitmap.

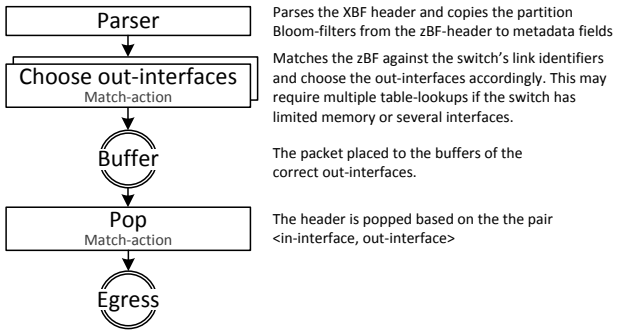
Parser
Parses the XBF header and copies the partition Bloom-filters from the zBF-header to metadata fields
Choose out-interfaces
Match-action
Matches the zBF against the switch's link identifiers and choose the out-interfaces accordingly. This may require multiple table-lookups if the switch has limited memory or several interfaces.
Buffer
The packet placed to the buffers of the correct out-interfaces.
Pop
Match-action
The header is popped based on the the pair <in-interface, out-interface>
Egress

Figure 12: P4 implementation of the popper and the relevant parts of the switch's processing pipeline.

The pipeline of the P4 popper-switch is illustrated in Figure 12. The pipeline starts with parsing the XBF packet and extracting the partition Bloom-filters from the zBF. The parsing is a simple operation done with a state machine. When the packet has been parsed, the out-interfaces are chosen based on the values in the zBF headers (i.e., the popper does not have to check the Bloom-filter in the IPv6 header at all). Choosing the correct out-interfaces may require several table-lookups if the switch has several interfaces and does not support large tables – a switch with 8 interfaces can choose the correct out-interfaces with one lookup on a table with 256 entries.

When the out-interfaces have been chosen and written to the packet metadata, the packet is replicated to the corresponding egress-interface queues. In the egress part of the pipeline, the switch performs a lookup based on the pair $\langle$*in-interface*, *out-interface*$\rangle$. As a result of this lookup, the packet is either popped or forwarded without modification to the egress queue.

We acknowledge, that the details for an efficient XBF implementation are highly depended on the target hardware. However, our P4 prototype shows that XBF can be implemented with relative ease to very simple devices. The only modification we had to implement on the P4 virtual-switch *bmv2*, which we used for testing, was to change the way how multicast groups are interpreted.

8. DISCUSSION AND FUTURE WORK

On the false-positive-free forwarding: The design of XBF, as described in this paper, requires that the fixed-length partition Bloom-filter have at least as many bits as there are links in the partition. Because of this design choice, XBF does not suffer from false positives. It, however, should be noted that XBF-like partitioning approach works also with conventional Bloom-filter based forwarding system that allows false positives. However, we believe that the false-positive-free

approach described in this paper is more palatable for two reasons.

First of all, probabilistically occurring false positives make network management and debugging very difficult. Second, by removing the possibility for false-positives, XBF solves the problems related to *forwarding anomalies* that have plagued Bloom-filter forwarding [15]. As an example, it is possible that a false positive takes the packet to an earlier point in the multicast tree. This kind of situation would effectively cause an infinite loop in the network. False-positive-free XBF does not suffer from these problems. There are several existing solutions to solve the problems caused by forwarding anomalies which, however, all increase the forwarding complexity. (We refer the reader to Särelä et al. for a more comprehensive description about forwarding anomalies [15]).

On the addressing: XBF is basically an addressing scheme where every link in the network is identified with a unique bit in a bitsring. It may appear couterintuitive that embedding this information into every packet is feasible. This, nevertheless, is the case as we have shown. The key insight is that the links in a multicast tree are not independent from each other. Thus, the network partitioning becomes an effective way to compress the information embedded into the packets.

In a more abstract sense, XBF can also be described as a traffic aware addressing scheme – if there is a lot of traffic between two nodes, they are more likely to be put in a single partition. We are not aware of any previous work that would have suggested optimising the addressing based on the network traffic patters.

Adapting to changing traffic patterns: In Section 3, we claim that the topology manager would perform the partitioning only once. However, traffic patterns may change over time thus rendering the partitioning ineffective. In fact, diurnal mobility patterns and changes in the traffic volumes are often seen in real networks [33]. Because of this, there should be some way how the partitioning can be changed without disrupting the ongoing communication – naive repartitioning would render the previously used Bloom-filters invalid causing extra delays.

The solution to this is to allow several coexisting partitionings. The idea is to use some header field, such as IPv6-header's traffic-class field, to identify a "partitioning" that a multicast tree uses. Obviously two multicast trees that use different partitioning cannot be combined. Nevertheless this mechanism could be used to reduce the overhead of the repartitioning. In addition to the signalling cost, the only drawback of this is that the flow-tables would grow n-fold where n is the number of coexisting partitionings.

Handling topology changes: Although Bloom-filter forwarding is intended for relatively static networks, no network is immune to occasional topology changes. In general, XBF handles changes in the network topology as any other Bloom-filter forwarding scheme. That is, if one or more links disappear from the network (e.g. due to a switch malfunction), these link-identifiers are no longer used when constructing paths. Moreover, when new links are introduced to the network, the topology manager simply assigns them new link-identifiers (this may require creating an entirely new partition). If this leads to non-optimal partitioning, the network can be repartitioned, as described above, once the topology has remained static for certain time.

Due to the fact that the topology changes can be handled as done in other Bloom-filter forwarding schemes, we refer the reader to [34] for more information about this matter.

Security: Bloom-filter forwarding has several well known security problems. In general, it should not used in untrusted environments without augmenting the protocol with hop-by-hop security checks that would significantly increase the forwarding complexity [21]. XBF, however, does not have this problem as it is intended for intra-domain use, where all nodes are trusted. Extending the work for inter-domain forwarding is left as a future exercise.

9. CONCLUSION

In-packet Bloom-filter forwarding was originally advocated as simple and fast way to do multicast forwarding. While this promise holds with small networks and multicast groups, the technology does not scale well – if the network or multicast groups become too large, the forwarding efficiency decreases due to false positives. The previously proposed fixes to this scalability problem have mainly relied on varying the length of the in-packet Bloom-filter. These solutions, however, increase the forwarding complexity significantly effectively countering one of the original design goals which was to enable simple and fast multicast forwarding.

The XBF-forwarding scheme proposed in this paper fills this gap. It makes Bloom-filter forwarding scalable by partitioning the network into equisized partition that reflect the network topology and traffic patterns, and uses a separate fixed-length Bloom-filter in each of these. Because of this, the forwarding inside every partition is kept very simple regardless of the topology or multicast-group size. Also, the partitioning also allows us to to eliminate the false-positives from Bloom-filter-based forwarding. Our evaluation shows that XBF is implementable with the existing SDN technologies and scales to very large networks.

10. REFERENCES

[1] P. Jokela, A. Zahemszky, C. Esteve Rothenberg, S. Arianfar, and P. Nikander, "Lipsin: Line speed

publish/subscribe inter-networking," *SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 4, pp. 195–206, Aug. 2009.

[2] J. Tapolcai, A. Gulyas, Z. Heszbergery, J. Biro, P. Babarczi, and D. Trossen, "Stateless multi-stage dissemination of information: Source routing revisited," in *Global Communications Conference (GLOBECOM), 2012 IEEE*, Dec 2012, pp. 2797–2802.

[3] A. Zahemszky, P. Jokela, M. Särelä, S. Ruponen, J. Kempf, and P. Nikander, "Mpss: Multiprotocol stateless switching," in *INFOCOM IEEE Conference on Computer Communications Workshops, 2010*. IEEE, 2010, pp. 1–6.

[4] R. Martinez-Aguilar and G. Fernandez, "Implementation of stateless routing mechanisms for multicast traffic on netfpga card," in *Communications and Computing (COLCOM), 2015 IEEE Colombian Conference on*, 2015, pp. 1–5.

[5] C. Macapuna, C. Rothenberg, and M. Magalhaes, "In-packet bloom filter based data center networking with distributed openflow controllers," in *GLOBECOM Workshops (GC Wkshps), 2010 IEEE*, Dec 2010, pp. 584–588.

[6] D. Li, Y. Li, J. Wu, S. Su, and J. Yu, "Esm: Efficient and scalable data center multicast routing," *Networking, IEEE/ACM Transactions on*, vol. 20, no. 3, pp. 944–955, 2012.

[7] N. Fotiou, P. Nikander, D. Trossen, and G. C. Polyzos, "Developing information networking further: From psirp to pursuit," in *Broadband Communications, Networks, and Systems*. Springer, 2012, pp. 1–13.

[8] M. AL-Naday, S. Hadjitheophanous, G. Petropoulos, M. Reed, J. Riihijarvi, S. Robitzsch, S. Spirou, D. Trossen, and G. Xylomenos, "POINT: IP Over ICN - the Better IP, Deliverable D3.1," Tech. Rep., December 2015. [Online]. Available: https://www.point-h2020.eu/wp-content/uploads/2015/12/POINT_D3.1-First_Platform_Design_v1.01.pdf

[9] D. Trossen, M. J. Reed, J. Riihijarvi, M. Georgiades, N. Fotiou, and G. Xylomenos, "IP over ICN – the Better IP?" in *Networks and Communications (EuCNC), 2015 European Conference on*. IEEE, 2015, pp. 413–417.

[10] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese *et al.*, "P4: Programming protocol-independent packet processors," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 87–95, 2014.

[11] S. Ratnasamy, A. Ermolinskiy, and S. Shenker, "Revisiting ip multicast," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 4, pp. 15–26, Aug. 2006.

[12] S. Rizvi, A. Zahemszky, and T. Aura, "Scaling bloom filter based multicast with hierarchical tree splitting," in *Communications (ICC), 2012 IEEE International Conference on*, June 2012, pp. 2790–2795.

[13] "isBF: Scalable in-packet bloom filter based multicast," *Computer Communications*, vol. 70, pp. 79 – 85, 2015.

[14] F. Hao, M. Kodialam, and T. V. Lakshman, "Building high accuracy bloom filters using partitioned hashing," in *Proceedings of the 2007 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '07, 2007, pp. 277–288.

[15] M. Sarela, C. Rothenberg, T. Aura, A. Zahemszky, P. Nikander, and J. Ott, "Forwarding anomalies in bloom filter-based multicast," in *INFOCOM, 2011 Proceedings IEEE*, April 2011, pp. 2399–2407.

[16] H. Lee and A. Nakao, "Improving bloom filter forwarding architectures," *Communications Letters, IEEE*, vol. 18, no. 10, pp. 1715–1718, Oct 2014.

[17] J. Tapolcai, J. Biro, P. Babarczi, A. Gulyas, Z. Heszberger, and D. Trossen, "Optimal false-positive-free bloom filter design for scalable multicast forwarding," *Networking, IEEE/ACM Transactions on*, vol. PP, no. 99, pp. 1–1, 2014.

[18] W. Yang, D. Trossen, and J. Tapolcai, "Scalable forwarding for information-centric networks," in *Communications (ICC), 2013 IEEE International Conference on*, June 2013, pp. 3639–3644.

[19] P. Mahadevan, D. Krioukov, M. Fomenkov, X. Dimitropoulos, k. c. claffy, and A. Vahdat, "The internet as-level topology: Three data sources and one definitive metric," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 17–26, Jan. 2006. [Online]. Available: http://doi.acm.org/10.1145/1111322.1111328

[20] S. P. Borgatti, "Centrality and network flow," *Social Networks*, vol. 27, no. 1, pp. 55 – 71, 2005.

[21] M. Antikainen, T. Aura, and M. Sarela, "Denial-of-service attacks in bloom-filter-based forwarding," *Networking, IEEE/ACM Transactions on*, vol. 22, no. 5, pp. 1463–1476, Oct 2014.

[22] F. Bourse, M. Lelarge, and M. Vojnovic, "Balanced graph edge partition," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2014, pp. 1456–1465.

[23] B. Hendrickson and T. G. Kolda, "Graph partitioning models for parallel computing," *Parallel computing*, vol. 26, no. 12, pp. 1519–1534, 2000.

[24] "Balanced graph partitioning," *Theory of Computing Systems*, vol. 39, no. 6, 2006.

[25] G. Karypis and V. Kumar, "A fast and high quality multilevel scheme for partitioning irregular graphs," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 359–392, 1998.

[26] G. Karypis, "Metis: A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices – version 5.1.0," Department of Computer Science, University of Minnesota, Tech. Rep., 2013.

[27] M. J. Reed, M. F. Al-Naday, N. Thomos, D. Trossen, G. Petropoulos, and S. Spirou, "Stateless multicast switching in software defined networks," *CoRR*, vol. abs/1511.06069, 2015.

[28] J. E. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin, "Powergraph: Distributed graph-parallel computation on natural graphs," in *Proceedings of the 10th USENIX Conference on Operating Systems Design and Implementation*, ser. OSDI'12, 2012, pp. 17–30.

[29] N. Spring, R. Mahajan, and D. Wetherall, "Measuring isp topologies with rocketfuel," in *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4. ACM, 2002, pp. 133–145.

[30] D. Vega, L. Cerda-Alabern, L. Navarro, and R. Meseguer, "Topology patterns of a community network: Guifi. net," in *Wireless and Mobile Computing, Networking and Communications (WiMob), 2012 IEEE 8th International Conference on*. IEEE, 2012, pp. 612–619.

[31] P. Elias, "Universal codeword sets and representations of the integers," *Information Theory, IEEE Transactions on*, vol. 21, no. 2, pp. 194–203, 1975.

[32] "P4 software switch, aka bmv2." [Online]. Available: https://github.com/p4lang/behavioral-model

[33] H. J. Fowler and W. E. Leland, "Local area network characteristics, with implications for broadband network congestion management," *Selected Areas in Communications, IEEE Journal on*, vol. 9, no. 7, pp. 1139–1149, 1991.

[34] A. Zahemszky and S. Arianfar, "Fast reroute for stateless multicast," in *Ultra Modern Telecommunications & Workshops, 2009. ICUMT'09. International Conference on*. IEEE, 2009, pp. 1–6.